Presented at the International Astronautics Congress, October 2004

PLANNING MARS MEMORY: LEARNING FROM THE MER MISSION

Charlotte Linde
NASA Ames Research Center, USA
Charlotte.Linde@nasa.gov

ABSTRACT

Knowledge management for space exploration is part of a multi-generational effort at recognizing, preserving, and transmitting learning. Each mission should build on the learning in the successes and failures of prior missions. Learning is the first step in knowledge production. The Mars Exploration Rover mission provides an opportunity to track how learning occurs, how it is recorded, and whether these representations might be optimized for subsequent missions. This paper focuses on the MER science and engineering team during rover operations. A NASA team conducted an observational study of the work and learning of this team. Learning occurred in a wide variety of areas: running two teams on Mars time for three months; using the instruments within the constraints of the martian environment, the deep space network, and mission requirements; planning science strategy; using the available software tools effectively. This learning is preserved in many ways. Primarily it resides in people's memories, to be carried on to the next mission. It is also encoded in stories, in programming sequences, in published reports, and in lessons learned activities. Studying learning and knowledge development as it happens allows us to suggest proactive ways to capture and use it across multiple missions and generations.

FULL TEXT

Recognizing and Preserving Learning

The idea of knowledge management assumes the existence of "knowledge". Yet knowledge is not a collection of natural objects. Rather, it is the result of learning by individuals, groups, and institutions. Thus, an important aspect of knowledge management, both for space applications, and for government and industry more generally, is learning how to recognize learning, preserve it, organize it, and make it available in a useful way.

Some of the knowledge captured in knowledge management efforts is represented in the form of documents generated as part of the ongoing process of a mission: design documents, records of mission reviews, records generated during the mission, etc. However, there is also a great deal of retrospective activity, attempts to recognize and preserve learning at the end of the project, or even more challenging, years later, from retiring experts. These post-project capture efforts most often consist of requirements to fill out computer-based forms in Lessons Learned systems. Somewhat less frequently, the lessons are produced by an interview process which results in written summaries or video records of stories told by soon-to-retire experts.

In order for such efforts to result in useful knowledge consistently and effectively, it is important to ask what learning is, and how we would recognize it if we saw it. While so-called Lessons Learned databases are common, their underlying definition of a lesson is frequently not clearly defined. Most often, a lesson is taken to be the product of a mistake or a mishap: something which must be documented to assure that it will not happen again. Many such "lessons" are to be found in mishap reports by investigation boards, and many reporting databases have as their primary purpose the collection of data to be provided to such a board in case of a mishap. This is a negative view of learning as an unpleasant activity, as expressed in the colloquial phrase: "Have you learned your lesson yet?"

Yet in spite of mishaps at worst and reinventions of the wheel at best, much new and positive learning **does** occur. As NASA prepares for a return to the moon after thirty years, we do know much more than we did then, both technically and institutionally. Thus, rather than merely focusing on lessons learned in the narrow sense of documenting mistakes, there is a prospective and more inspiring question: how do people, projects, and institutions learn new things, how is this learning preserved and continued, and how can we plan for such learning and preservation across a multigenerational mission?

There are many types of learning in projects as large, complex, and lengthy as space exploration missions. At the highest level, there is mission-scale learning. This includes such issues how to maintain support for a mission over many years, multiple elections, and changes of government commitments to the mission, and how to manage across changes in the requirements for how a government agency is legally mandated to work with contractors, and in the mandated system for management processes. (For example, (1) describes changes in systems engineering models in American and European space programs.) At perhaps the most detailed level, there is learning about individual components and their behavior. The behavior of the O ring at low temperatures is perhaps the most notorious example, but any spacecraft design incorporates hundreds of thousands of such lessons learned about components and their interactions. At an intermediate level, there is the learning generated in the design process: inheritance of previous mission design concepts; development of new designs, materials, and processes; discovery by iterative design, etc. This is true both for hardware and software design. But it is perhaps most extreme in the design of software, since arguably, this is the area in which the greatest changes in the state of the art are developed from mission to mission.

In order to describe what learning looks like as it happens, this paper focuses on an intermediate level of learning that falls between the level of government policy and the level of widgets, looking at ongoing learning during mission operations, and the ways in which this learning is formally or informally incorporated into the memory of the institution performing the mission. It is difficult to capture this kind of learning without special efforts. Such learning may become part of the mandated or informal work procedures, incorporated into software code, or may simply become what "everybody knows." By the end of a mission, it is easy to forget that learning has taken place, since it has become obvious "common sense." Some learning may be captured but stored in disparate places that would be very difficult to reconstitute. Finally, without the intention and ability to recognize this kind of learning on the fly as learning, there is no reason to try to capture it in a systematic way.

This paper first describes what such learning looks like, describes how such learning is currently captured (or not), and finally, discusses methods for discovering, capturing, and representing these more subtle forms of learning. In effect, the paper describes a few examples of the many learning processes which produce the data and information which knowledge management systems attempt to manage.

Research Site and Methods

This paper uses as data the current Mars Exploration Rover mission: astonishingly current in fact. As this paper is being written, the mission is still ongoing, approximately 5 months after its nominal duration.

As probably every reader of this paper already knows, NASA's Mars Exploration Rover Project consists of two rovers performing robotic geological fieldwork on two locations on the surface of Mars, searching for evidence of a history of past water activity. Sequences of launch, cruise, and arrival operations dispatched each rover to a different area of the planet three weeks apart to explore those areas for about three months each. The first rover, Spirit, landed on Mars on January 4, 2004, followed by the landing of the second rover, Opportunity, on January 25, 2004.

The rovers were designed to recognize and maneuver around small obstacles on their way to target rocks selected by scientists from images sent by the rovers. They carry an instrument package which includes a panoramic camera (Pancam), a stereo camera with 11 color filters and 2 filters for imaging the Sun; a miniature Thermal Emission Spectrometer (Mini-TES); a robotic arm with a microscopic imager to produce extreme close-up images of rocks, soils and particles; an alpha-particle spectrometer to detect elemental abundances of rocks and soils; a Mössbauer spectrometer to distinguish iron-bearing minerals; and a Rock Abrasion Tool (RAT) for brushing and grinding. Additionally, the rovers contain two monochromatic navigation cameras, four hazard avoidance cameras, and high gain and low gain antennae for data transfer via Deep Space Network (DSN) or Direct to Earth (DTE) transmission.

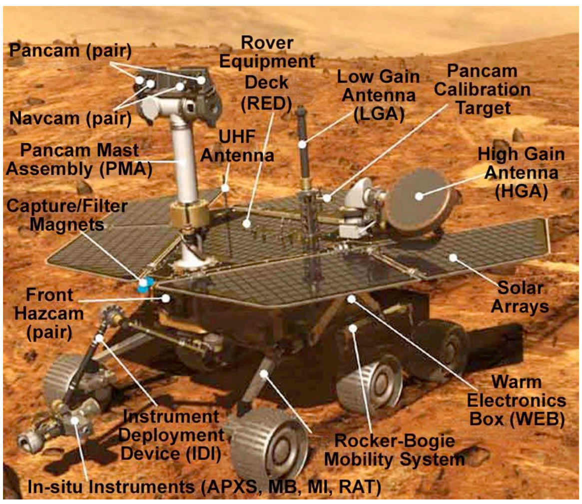

The Rover's Instruments

In addition to providing unprecedented data about the planet, the MER mission also has provided us with an opportunity to track how the science and engineering team learned, how this learning was recorded, or assimilated without specifically being marked as learning, and whether the representations of this learning are likely to be usefully available for subsequent missions. This paper is based on an observational study of the work of the MER science and engineering teams. Conducted by members of NASA ARC's Work System Design and Evaluation team, the study includes ethnographic observation and video and audio recording of the work of the team, analysis of documents, and interviews with team members. Additionally, as a result of NASA-wide attempts to capture Lessons Learned in preparation for the planning of NASA's Vision for Space Exploration, knowledge capture workshops for Lessons Learned have been held at all the NASA centers, and the JPL workshop includes managers from the MER project, thus furnishing additional data.

Grossly simplified, the work of the science operations team consists of receiving downlinked data, assembling it into usable data products, analyzing the data to determine what data to request next, producing plans for the rover's activities on the next sol (martian day), and creating command sequences to uplink to the rover.

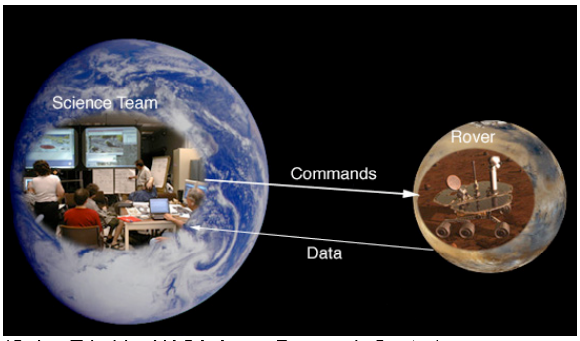

(© Jay Trimble, NASA Ames Research Center)

Rover Operations Communications Model

The science team is divided into five theme groups by discipline: Atmospheric Sciences, Geology, Soils and Physical Properties, Geochemistry/Mineralogy, and Long Term Planning. Long Term Planning is a special group whose charge is to develop a general plan for each sol, and ensure that there is continuity between sols, and that science and mission objectives are tracked and accomplished. Discussion, learning, negotiation, and planning occur in a series of formal meetings, punctuated by individual or *ad hoc* small group work.

Types of Learning: Science and Operations Team

There were many forms of learning and learning capture observed during the work of the science and operations team. Indeed, this paper describes a work in progress, since at the time of writing the mission still continues, months longer than the nominal 90 day duration of each rover's mission. Further, although one issue of *Science* has already appeared with initial findings about the first 90 days of the Spirit Rover, much of the work of scientific analysis and assessment of design and operations decisions will be performed only after the conclusion of the mission. However, even in the heat of the moment, it is possible to sketch some of the major areas of learning and knowledge preservation.

Learning to Understand the Instruments

A major type of learning we observed was the team coming to understand the potentials and limits of the instruments of the rover, and how to work with them to obtain the science return desired. The team had to strike a balance continuously between the scientific desire for as many observations as possible, made at the highest resolution possible, and the practical restrictions imposed by available power on the rovers and bandwidth for downlink on the Deep Space Network, which serves the data transmission requirements of many missions simultaneously.

Let us consider in detail one example of many, observed three weeks into the mission. In an informal meeting, a group of scientists discussed possibilities of compressing data from the Microscopic Imager to overcome a potential problem of too high a data volume if observations are made without compression. The options are either to reduce the number of bits per pixel or to take fewer observations, or some combination of these strategies. The discussion includes expertise offered by the writer of the data compression algorithm, who says that "1 bit per pixel is far from adequate' while 3 bits cannot be distinguished from lossless compression." One of the scientists showed the results of a rapid experiment he has performed in the last 20 minutes, producing images at a low pixel rate. The chair of the Long Term Planning Theme Group suggested a compromise, and an immediate decision was made that a rate of 3 bits per pixel with 5 rather than 7 slices would adequate for the current task, but that the decision might be revisited for observations on different soil.

This is an example of practical learning on the spot: learning how to use an instrument in new conditions within the changing constraints of the situation. It becomes an immediate part of what the team knows: a certain compression rate for this instrument will be adequate for certain scientific purposes, and thus can relieve pressure on the problem of the volume of data to be downlinked.

How is such learning captured or represented? Consider the following quote from a report on the Microscopic Imager in Science (d)

"A typical MI data set includes a stack of three, five, or seven MI images, acquired at 3-mm steps along the MI optical axis with the dust cover open. This acquisition approach helped to ensure optimum focus on targets with relief greater the MI depth of field. The number of images in an MI stack was kept small to minimize the volume of extraneous MI data returned to earth. Most of the MI stacks included at least one image in good focus, but uncertainties in the front Hazcam terrain model resulted in poorly focused MI images in some cases. Color information was sometimes added by acquiring an additional single-frame MI image of the same target at the nominal best-focus position with the dust cover closed. Some of the MI targets were imaged with a binocular stereo pair of left/right MI stacks or even a mosaic of MI stacks."

This dense paragraph reports tersely the result of hundreds of discoveries of how best to use the instrument, each of which was the result of discussions like the example above.

Indeed, as the mission continues with degradation of the rovers' condition and progressive loss of battery power, further learning is taking place about how to use the instruments under extreme conditions, including maneuvers which would not have been even attempted during the nominal duration of the mission, for fear of damaging the rovers.* Without the observation of the actual work of the team, it would be hard to retrieve from this archival record the details of the learning which produced it.

Learning to Understand the Limits of the Models

A major part of the daily tactical planning process involved scientists individually, in small groups, and in the formal science team meeting of the day, using a software tool called Science Activity Planner (SAP). This tool allows the scientists to specify the observations needed to produce a required observation (and its resulting data product), and provides estimates of the amount of time and power that any given observation would require. It thus allowed scientists to decide on a preliminary version of the particular activities that would be requested on a given day.

Thus, for example, if the scientists were considering a drive of a given number of meters, followed by pictures to be taken by specified cameras, SAP could show an estimate of the time it would take to drive to the spot, as well as the amount of power available after the drive. However, these estimates were quite rough. Over time, scientists learned that they could produce plans with power requirements well over the 100% available power figure indicated by SAP.

Similarly, odometer figures for the rover were quite inadequate, and required compensatory calculations for where the rover actually was located, as opposed to the location indicated by the odometry data.

For example, in one argument about the strategy for the coming sol, one scientist reminded the group of previous problems with odometry which might compromise the current plan:

> "We may have a situation like we did with Faux Trench that our pointing is compromised by the terrain in front. We drove away from Faux Trench with the hope of turning around and shooting it. And we missed **wildly** because the odometry apparently was wildly wrong. And we may be in a condition like that again. So my point here is that MiniTES on those rocks will probably be a bit dicier than it has been in previous times."*

This example shows the group in the process of learning what activities are possible: a proposed plan is critiqued because the requested images require knowing the precise location of the rover, which can not be guaranteed given the inaccuracy of the odometry in the particular terrain. Some of the problem of calibrating the odometry has already been discussed in the first publications of the Athena Science Team. For example, a discussion of how the exact path of the rover was determined includes the note that "Results indicate that the total length of the rover traverse during the 90-sol primary mission was 637 m as measured by wheel odometry, 506 m as assessed from formal localization performed for the locations where image data were acquired." (3)

In general, experience allowed the scientists to move from the more conservative estimates provided by the models to more and more generous estimates in their planning. This kind of learning, however, raises a real question for knowledge management. We do not know whether the SAP software will be carried over for use in future rover missions. If it is, we must ask what mechanisms should be in place to collect the learning about the relative accuracy of the model's predictions, so that the next version can benefit from this learning. The issue is even more difficult if another modeling program is developed: how can this kind of knowledge of the rover's actual behavior under various conditions be preserved to be incorporated into the

* At the time of writing, August 29, 2004, both rovers were will in operation, as much as 5 months past the nominal completion of the mission.

* This quote is taken from the Science Assessment Meeting, Spirit Sol 66, March 10, 2004.

development of the next generation of models?

Strategic Planning

One important development, which made the mission so fruitful, was the science and engineering teams learning how to do long term strategic planning, rather than operating only sol by sol. Obviously, the mission had overall strategic goals, both scientific goals ("follow the water"), and mission goals (90 sols of operation, 600 meters driven, etc.). However, the challenge for the team was to bridge between these very high level goals and the intense time pressure of operating during any given sol. Initially, during the field tests which served as training both in how to run the rover and how to perform mission operations, scientists tended to plan one sol at a time. However, as they grew more familiar with the planning process, even in the field tests and certainly during the actual mission, they became increasingly adroit at producing branching multi-sol and even multi-week plans. Strategic planning involved learning how to negotiate a number of tensions : the possibly competing demands of different science theme groups, the "bird in the hand' problem", and the development of strategies for producing multi-sol plans with software designed for single-sol planning.

Competition and collaboration among theme groups: The initial model during the field tests was of the different theme groups working separately to develop their plan for the next sol and then negotiating what would actually happen during meetings of all the groups. However, as the theme groups grew more familiar with the instruments and the process, they began informal collaborations immediately, working both with intellectually adjoining groups, and with the long term planning group, which served as the custodian of strategic aims. This allowed for the development of branching plans that ensured that the differing data requirements of the theme groups would be fulfilled over several sols, if not in a given sol.

Additionally, an informal process of monitoring fairness developed. Members of all the Theme Groups noticed situations in which the observations requested by a given Theme Group had not been performed for several sols, and attempted to ensure that those observations were made within a reasonable time. (This was a particular issue for the Atmosphere Theme Group, since its desired observations of sky or horizon were not as dependent on specific locations, and hence on performance on a particular sol, as were those of other theme groups, which required observations of specific rocks or soils.) Interestingly, there were occasions when members of one theme group pointed out that another group's requested observations had been postponed for many sols, and that fairness required their requests be given priority.

"Bird in the hand problem" : There was a constant tension for scientists between taking as many observations as possible of the location they were in, as opposed to moving on to more obviously interesting locations. The problem, obviously, was that while some locations were more likely to be more geologically significant than others, so little is known about Mars that every location has its charms and possibilities. The tension was heightened by the fact that the rovers had a limited but unknown duration of operation. It became the task of the chair of the science team (a position which rotated daily or nearly daily) to arbitrate these arguments.

Strategies for multi-sol planning: one problem which the team faced was that the computer tools for science activity planning were designed to handle one sol at a time. It was not possible to carry over requested observations automatically : such bookkeeping was handled by human memory. Such multisol planning, as well as the tension between strategic and tactical aims, was assisted by several developments. One was a program for representing branching possibilities across many sols. This was the Sol Tree, a representation of branching possibilities for multiple sols, developed by the Long Term Planning Group, and usually displayed publicly using the large plasma screen monitor of a collaborative work tool called the MERBoard

Another way of handling the complexities of planning was the development of standard strategies and names for them. This is a well known cognitive strategy : chunking – grouping objects or actions into a newly defined entity, and lexicalization – development of new names for these entities. An enormous amount of learning is encoded in these names, and provide a likely mechanism for the preservation of this knowledge across missions. Some examples follow.

At the highest level are names describing multi-sol strategies which have become conventionalized. For example, late in the mission, the team developed a 'lily pad strategy' : a plan for moving the rover from one to another north-facing location in order to take advantage of available solar power, as the advent of the martian winter brings shorter daylight periods, with the Sun in the north.

At the next level down are names for sol types. For example, the 'touch-and-go' sol, which was featured in description of the mission in the first special issue of *Science* devoted to Spirit at Gusev Crater. (2) "Spirit drove about 600 meters from the Columbia Memorial Station to the rim of Bonneville Crater. Along the way, the rover performed quick analyses, called touch-and-go operations, in which the instruments on the arm touched and analyzed a feature but no brushing or grinding was done. … Exploration has changed since the era of nautical exploration by large sailing vessels, when 'touch and go' is thought to have originated to describe a ship's keel touching the seafloor briefly but not getting stuck."

At a lower level is the combination of several activities, which may be performed during a single sol. These include the 'driveby' – driving past a specific target and photographing it ; the 'scoot and shoot' – a series of drivebys ; a 'scratch and sniff' – drilling into a rock and then examining it with the rover's other instruments ; the 'stutter stop' – stopping a drive one meter before the planned end, taking images, and then continuing to the destination. Additionally, the team developed conventional terms for types of images, types of features, data downlink issues, and standard strategies for using the science planning software.

Learning How to Learn

All of these are examples not only of learning the specific technical strategies and solutions, but also of learning how to learn as a team, under strict constraints of time, available power, transmissible data volume, and negotiated group demands and strategies. The examples cited here show the learning happening in small groups, taking advantage of the co-presence of team members. That is, the experts needed for consultation were either present, or available by phone (since presumably the team members working on a given rover were operating on the sleep schedule of that rover's time zone.) Team members thus can know who the experts are on particular instruments, software, etc., and can bring them in smoothly to contribute their expertise.

Additionally, members of the science and engineering teams moved between rovers, depending on personal schedule issues and scientific interests. Such movement allowed for learning to be transmitted almost immediately from one rover team to another.

This poses an important question for the design of mission operations for future missions. For example, MSL '09 has an expected duration of three years, which clearly precludes operation on Mars time, and which would not allow participants from institutions other than JPL to be present at JPL for the entire course of the mission. Planning for this kind of learning, and for its dissemination across multiple institutions, nations, and continents requires first recognizing the value of these informal, on-the-fly collaborations, and then working seriously on the design of technology for remote collaboration that can allow for analogues or substitutes.

How Are These Types of Learning Preserved (or Lost)

Thus far, this paper has attempted to sketch some of the many types of learning which happen as an unremarked part of the daily work of a mission. Obviously, the work of the MER science and engineering teams is anything but routine: team members are creatively performing new and startlingly complex and successful activities in an unexplored terrain on another planet. Certain examples of the team's learning were obvious to the entire world: the software problem with flash memory and the successful fix for it was the stuff of news headlines and ongoing drama. But most of the learning discussed here is not dramatic in this sense. Rather it is daily, ongoing, incremental. The question for knowledge management efforts is how to recognize this *as* learning, and ensure that it is preserved.

If knowledge management begins to view knowledge as the product of learning, it is necessary to ask: 'Who learns?' Certainly, individuals learn. Does it also make sense to say that institutions learn, or that the scientific community as a whole learns?

Individuals learn, and as they move from project to project or mission to mission, there is a good chance that their learning will be available on the next project. Individuals also train and mentor their colleagues and their juniors, providing another way that individual learning is transmitted.

But what would it mean for an institution to learn? Institutions learn by retaining people who have learned, and by engaging them to train others. In terms of data, in a weak sense, it could be argued that an institution has learned by collecting relevant data and making it available as knowledge in archives, databases, etc. However, this constitutes actual learning only if the knowledge is both usable and used. In the strongest sense, an institution has learned when its behavior has been changed. Learning has happened successfully when it has been incorporated into the procedures mandated by the institution, into the design decisions the institution makes, and into the ways in which the institution functions internally and with external partners.

Let us now consider strategies for ensuring that learning is preserved as knowledge, at the individual, team, institutional and scientific levels.

Preserving Individual Learning

Although this section discusses individual learning, it must be stressed that in fact no learning can be purely individual, with no contribution to or from other people and groups. This may appear to be a tendentious claim in the case of a mathematics student sitting alone in a room with a textbook (but how was the textbook produced?). In the case of the MER mission, though, the work is so massively collaborative, depending on the contributions of so many different participants, that learning too must be collaborative.

At the same time, particular people do learn, and certain types of learning are carried on only by people, with little or no contribution from documentation. There are a number of ways for institutions to preserve this learning: career track planning, mentoring, succession planning, etc.

The career tracks of individuals can make a large difference in whether their learning is preserved for the larger institution. Presumably many of the JPL employees who worked on the MER mission will move on to later Mars missions, thus ensuring that what they have learned is available for these missions. Opportunities for mentoring were present during the MER mission in a number of ways. Senior team members, who had worked on the Viking and Pathfinder missions, were part of the team, many of them in roles that gave them responsibility for the strategic planning of the mission. Many of the younger science team members were graduate students of the academic co-investigators: again, a form of mentoring built into the structure of graduate education. They will be well positioned to propose to become co-investigators or principal investigators for future missions. However, in retrospective assessment of the operation of the mission, one question worth asking would be whether there were sufficient opportunities and time for effective mentoring to take place at every level of the team structure.

Some individuals take steps to preserve their individual learning by publishing accounts of their experiences (e), (f). However this is a strategy that appears to be available only to fairly high level scientists and managers, and to accounts of successful missions.

Preserving Group Learning

One of the most important types of learning that happened was that the many people on the MER team, coming from many different institutions, disciplines, and levels of experience, learned to work together under difficult conditions as an extraordinarily effective team. It is impossible that the entire team could be brought together to work on the next major Mars mission, for example, MSL '09. However, some key people certainly will move to this mission and will bring with them some of this learning about effective teaming.

A particularly impressive example of this type of learning is the way the various Science Theme Groups and engineering groups learned to work together, allowing the planning process to occur earlier and earlier in the sol. A question for a retrospective assessment of the operations of the mission would be to determine how well the various groups understood the work of other groups and the constraints on it, particularly groups which worked at different times in the day, and so did not have the opportunity to learn through face-to-face communications.

The work of learning how to operate on Mars time provides a very detailed example of group learning. Veterans of the Mars Pathfinder mission recalled the difficulties of working on Mars time, and asked Human Factors experts at NASA Ames Research Center for assistance in mitigating the difficulties. A team began with a survey of Pathfinder personnel to determine what factors of Mars time operations were most difficult for them. Schedules, housing facilities, and policies were devised to make this difficult operation as easy as possible. This is learning deep down in the details. Mundane but critical examples include providing short-term rental apartments for visiting scientists located in quiet neighborhoods near all-night markets, learning how to convince the rental company to provide adequate blackout curtains to allow sleep during the day, providing on-site rooms with cots for "power naps", developing a JPL policy for reimbursing taxi expenses for team members who were too fatigued to drive home safely.

<u>Preserving Institutional Learning</u>

Institutions learn when their members have learned. Thus, developing effective training is an important part of any institutional learning strategy. However, in a more formal sense, institutions have learned when they have developed policies, procedures and practices which incorporate the learning developed by individuals and groups. This is a method which puts the responsibility for incorporating learning on the institution directly. As shown by the Columbia Accident Investigation Board Report (8), it is not enough to mandate procedures. It is also necessary to have continued organizational attention to ensure that these procedures are followed, and that they *can* be followed. The existence of conventionalized workarounds to institutional policies is not, as management might be tempted to think, a testimony to human perversity. Rather, it suggests the possibility that the institutional procedures themselves may be impeding the workflow required, or that additional, unacknowledged factors in the environment, such as production pressures, may make them difficult or impossible to follow. Additionally, the rapid tempo of a mission like MER suggests a need for attention to the meta-process of producing procedures. Is it possible to include new learning into the mandated processes during the mission, or does the validation process for processes introduce too long a delay?

<u>Scientific and Public Learning</u>

It is also important to note that knowledge produced by such a mission travels far beyond the people and the institutions which carry them out. Obviously, the MER mission has already resulted in scientific publications, and the data of the mission will be a source of scientific analyses and publications for decades to come. The mission was unusual in choosing to make its raw data publicly available on its website as soon as possible. It also made determined and intensive efforts at reaching the general public, with scientists working in a sustained way with press officers and science education experts to get relevant information out rapidly and continuously.

Conclusion: How Best To Preserve Knowledge

Let us conclude by considering how to implement strategies and tools for the preservation of knowledge at all these levels. First, this paper has argued that knowledge management should begin with an understanding that knowledge is the product of learning, individual and collective. Such a view provides a very different perspective for knowledge management than the common assumption that the beginning point for knowledge management is data, which is turned into interpreted information, and finally into usable knowledge. Taking this view suggests an emphasis on identifying where learning is happening, and determining the best ways to support and extend it within an institution, as well as preserving it for future uses.

Lessons Learned Activities

A common strategy to achieve institutional learning is to develop Lessons Learned databases for particular missions or projects. The most common way to do this, which forms one of NASA's Lessons Learned activities, is to set up a Lessons Learned database to which project participants are requested to contribute relevant lessons, which are then validated before entry into the official database. These Lessons Learned are available both publicly, and on an intranet site (llis.nasa.gov). In general, a search of lessons for both Pathfinder and MER missions shows a predominance of small scale technical issues. Additionally, the format of the Lessons Learned form requires a description of a driving event, a description of the lesson learned, and recommendations. This format is slanted towards problems rather than towards descriptions of positive learning. Further, the form is short, which precludes a detailed description of the kinds of learning and ways of learning described above.

In addition to the individual process of gathering Lessons Learned described above, there are also Lessons Learned workshops. These vary widely in format, and the variations make a great difference in the likely effectiveness of the product. The optimal arrangement is to have the activity as part of a scheduled post-mission review, and ensure that the participants include not only mission scientists and engineers, but also institutional decision makers who can derive immediate action items for their departments, and junior level employees, who can use the event as an opportunity to learn in preparation for later stages on their career tracks. Using a workshop as a way to populate a database is much less effective, although probably still better than having individuals fill out in the privacy of their own offices, since at least some discussion, argument and cross-learning is possible.

There have been a number of efforts in the past, as well as ongoing efforts for NASA's new Exploration Mission to extract and record knowledge from senior or retiring experts. These activities usually take the form of guided interviews, with the product being transcripts, edited versions or video tapes of conversations. These efforts share the problem of all oral history projects: the difficulty of effective indexing and formatting. Without a great deal of work after the knowledge capture phase, what is produced is raw data, in a form likely to be unusable for anyone who wants to come up to speed fast on a technical area. A number of technological solutions are being attempted in the area of automatic indexing of video records, and these may alleviate some of the problems. However, thus far, research has shown (9) that the most successful capture efforts of oral lessons learned involve a great deal of skilled editing to make them useable. Two outstanding efforts, the Center for Army Lessons Learned, and the Aviation Safety Reporting System, use highly skilled senior experts as interviewers and editors. Thus knowledge management has a lesson to learn from Hollywood: it's the postproduction rather than the filming that makes or breaks the product.

Preparing to Capture Future Learning

While producing Lessons Learned databases, and capturing departing expertise are valuable methods, which need refinement to increase their value, it is important for the field of knowledge management to develop other ways to accomplish knowledge capture as well. One suggestion would be for every major project to have a learning office or a learning officer, whose job it is to record successful learning as it happens. The MER mission did have a position called "documentarian" whose job was to record the events of each day. However, this position tended to be filled by graduate students, whose lack of experience led them to record tactical rather than strategic decisions.

Let us reflect on what we might want to learn from the Apollo lunar landing missions.

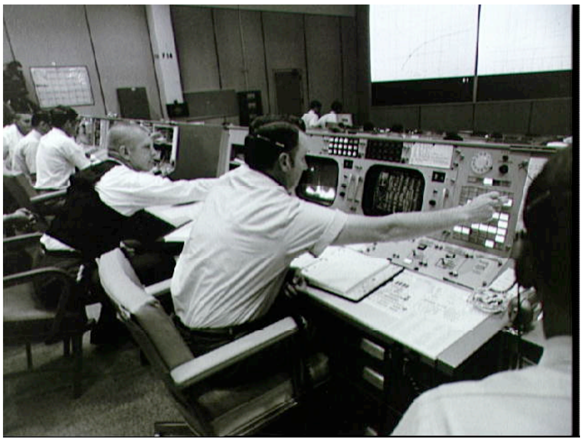

Apollo 11: Flight Director's console in the Mission Control Room during the launch of the Apollo 11 lunar landing mission.

While a return to the moon will use very different technologies and procedures, there is still much that was learned during those lunar missions that is still valid and important 35 years later.

While Lessons Learned databases, and capture of departing expertise are valuable methods, which need refinement to increase their value, it is important for the field of knowledge management to be looking for other ways to accomplish knowledge capture as well. As we prepare for a multi-decade, multi-generational exploration effort of the Moon and Mars, we must make every possible effort to make the best possible use and record of the learning we accomplish along the way.

Acknowledgements

I would like to thank my colleagues on the MER fieldwork team. In particular, I am grateful to Roxana Wales, who did the intensive, difficult, and usually invisible work of arranging for such fieldwork to be possible. I also thank my other team-mates Valerie Shalin, Chin Seah, Zara Mirmalek, and William J. Clancey. I am indebted to Cheri Brownton, Lisa Faithorn and Robin Kornman for help at various stages in the preparation of this paper. Particular thanks to all the members of the MER mission who generously allowed us to share this adventure with them.

References

(1) Johnson, Stephen B. *The Secret of Apollo: Systems Management in American and European Space Programs*, Johns Hopkins Press, (2002.)

(2) Rowan, Linda. 'Touch and Go', *Science*, **205**, (793) (2004).

(3) Squyres, S.W. et al "The Spirit Rover's Athena Science Investigation at Gusev Crater, Mars", *Science*, **205**, (793) (2004).

(4) Arvidson, R.E. et al. "Localization and Physical Properties Experiments Conducted by Spirit at Gusev Crater" *Science*, **205**, (821) (2004).

(5) Kerkenhoff, K.E. et al. *Science*, **205**, (825) (2004).

(6) Mishkin, Andrew. *Sojourner: An Insider's View of the Mars Pathfinder Mission*. Berkeley Books, (2003).

(7) Shirley, Donna and Danelle Morton. *Managing Martians* Bantam Dell, (1998).

(8) Gehman, Harold W. et al. *Columbia Accident Investigation Board Report*. NASA and Government Printing Office, (2003).

(9) Linde, Charlotte "Narrative and Tacit Knowledge Exchange." *Journal of Knowledge Management*, **5:2 (**160-170), (2001).